CCDE Practical Studies

Practice Lab Exam 1 – "iSize"

Full Practice Lab Examination & Debrief

Martin J Duggan CCDE 2016::6, CCIE#7942

About the Author

Martin J Duggan CCDE™ 2016::6 and CCIE™ #7942 is a Senior Network Architect and Cisco Press author. Martin gained his CCIE™ in 2001 and has been passionate about Cisco qualifications and mentoring ever since. Martin successfully passed his CCDE™ practical in 2016.

Previous Publications include:

Author of Cisco CCIE Routing and Switching v5.0 Configuration Practice Labs ISBN-10: 0-13-378631-5 http://www.ciscopress.com/store/cisco-ccie-routing-and-switching-v5.0-configuration-9780133786316

Author of Cisco CCIE Routing and Switching v5.0 Troubleshooting Configuration Practice Labs ISBN-10: 0-13-378633-1 http://www.ciscopress.com/store/cisco-ccie-routing-and-switching-v5.0-troubleshooting-9780133786330

Co Author of CCIE Routing and Switching Practical Studies ISBN: 1587051478 (Ciscopress) http://www.ciscopress.com/title/1587051478

Technical Editor of Cisco Field Manual, Catalyst Switch Configuration ISBN:1587050439 (Ciscopress) http://www.ciscopress.com/title/1587050439

Want to know more about Martin's journey to CCDE™?

https://learningnetwork.cisco.com/blogs/unleashing-ccde/2016/06/17/60-limit-by-martin-duggan

Want to know more about Martin's CCIE book?

http://www.gocertify.com/articles/interview-martin-duggan-author-of-ccie-routing-and-switching-v5-0-configuration-practice-labs.html

Technical Editors

Daniel Dib, CCIE™ #37149, CCDE™ #20160011, is a Senior Network Architect at Conscia Netsafe. He works with creating scalable, modular and highly available network designs that meet business needs. Daniel started out in implementation and operations and got his CCIE in 2012. In May 2016 he became the second person in Sweden to get CCDE™ certified.
He often acts as a subject matter expert for his customers with deep expertise in routing, switching, multicast and fast convergence.

Nicholas (Nick) Russo, CCIE™ #42518, holds active CCIE certifications in both Routing and Switching and Service Provider. Nick served 6 years in the US Marine Corps, many of which were in a technical networking capacity, then went on to become a professional network engineer as a civilian. Nick also holds a Bachelor's of Science in Computer Science, and a minor in International Relations, from the Rochester Institute of Technology (RIT). Nick lives in Maryland, USA with his wife, Carla, and their daughter, Olivia.

About the Book

This book is the first part of a 3 part series. I intended to release a single publication which included 3 complete CCDE practice lab exams and full debriefs for each but quickly realized that I could help CCDE candidates approaching their lab examinations if I released a lab at a time. Once all 3 lab exams have been completed I will release a bundled publication and include additional hints, tips and background information on the CCDE exam. These practice labs are the culmination of my journey towards CCDE and the thousands of hours of study I undertook in order to be successful.

Taking the CCDE lab exam was a real challenge for me, personally I found it significantly harder than achieving my CCIE back in 2001. The certification has been running for some 8 years at time of writing and there are only approximately 350 certified individuals to date. My biggest issue in joining the 349 others was practice, I just couldn't get enough quality practice exams that would prepare me for the technical marathon that is the CCDE lab exam. That's why I wrote this lab exam and why I will create an additional 2 further practice labs to help you achieve your goal of joining the CCDE club.

The book is in two parts, Part 1 covers the practice lab exam and Part 2 is the debrief to the exam. Don't even think about reading the debrief section until you have taken the practice lab.

You need to treat this practice lab as if you were taking the real exam, get somewhere quiet and comfortable where you can focus and imagine you are taking your exam for real. Follow the rules of the real lab exam and do not go backwards or forwards on the questions but feel free to read the background information and any emails that come in as much as you want to. Have a piece of paper to make notes and create drawings on as some of the questions will require diagrams and if possible highlight relevant information within documents to help you detail the facts which are ultimately requirements or constraints.

Read the background information and make notes but do so quickly, you should save as much time as possible for the questions. If you feel you don't have sufficient information to answer a question then head back to the documentation or emails as the information you need is there.

Don't worry about stopping at 2 hours and seeing how you have done, carry on until completion but do make a note of your time, if you only managed to get half way through by following the status bar included in the questions and you are at 2 hours just make a note that you need to improve your speed, significantly!

This practice lab is about as close as you can get to taking the real exam, I hope you enjoy it.

Part 1: Practice Lab

iSize

CCDE Lab Practice Exam 1

From: Sam Dukes
To: CCDE Candidate
Subject: Welcome on board

Hi

Welcome to iSize, we're glad to have somebody with your experience for the next 3 months assisting us with our current issues and to develop our network as we grow our business, exciting times ahead and a great opportunity for you.

Please take a look through the attached documentation to familiarise yourself with our company and network.

Rgds

Sam Dukes
CTO iSize

Document 1 – Company Overview

iSize is a bespoke bicycle manufacturer which has been trading for fifteen years and has seen rapid expansion in the previous three years. The company is located in Southampton, UK. It was formed in 2008 by two friends formally employed by IBM who enjoyed cycling together and frequently discussed the limitations of mass produced cycles. They found a gap in the market which wasn't being addressed by global cycling manufacturers. The major manufacturers address the common requirements for customers by supplying standard solutions in the on and off road markets in the form of multiple styles, sizes and budget ranges for popular cycle models. The founders of iSize determined that there would be a niche market for made to measure bicycles which were highly customisable and that clients would be prepared to invest in a bicycle which fitted them perfectly and suited their exact needs.

iSize grew rapidly after patents were successfully filed for their innovative designs and investment began to flow into the company. Their IT needs grew significantly which resulted in the hiring of additional staff specializing in the compute and application development environments. The network however was treated as a utility and suffered from lack of investment and personnel to support it, individual contractors were commissioned on an ad hoc basis as and when changes were required.

Currently when a new bicycle prototype is required it generally involves a site visit to Bi-kea (the manufacturing company in Taichung, Taiwan which has exclusive rights to produce iSize's product range), delay incurred due to creation of bespoke tooling and production run scheduling negotiation. iSize feel this model is not conducive to cost efficiency and delays time to market for their products.

iSize are investigating the use of prototype production in house by use of 3D printing using the latest composite printing technology, this would negate the travel requirement and allow them to provide virtually identical prototypes to factory manufactured cycles without the delay of the fabrication of tooling and booking of production runs within the factory. Longer term, iSize would like to wholly manufacture locally in the UK using the 3D printing model if possible to reduce the carbon footprint involved in manufacture and shipping of the bicycle frames from the far East but are acutely aware that this will not be achievable until their IT systems and 3D printers have scaled sufficiently allowing them to accomplish this.

In terms of IT, iSize have persistently focused on the compute environment increasingly using virtualization to increase utilization of their servers and have not modernized their network. Having used multiple contract resources to design and build the network they are acutely aware that the organic growth has introduced day to day issues which are ultimately attributed to the actual underlying network design. Network support is outsourced to a 3rd party company who are based 80Km from the HQ who normally replace devices when there are issues, typically a different engineer attends site each time an issue is reported and outages are becoming more common with increasing amounts of downtime. The CTO has stated that he is willing to invest in the network to improve the stability, supportability and scalability but that he hasn't secured a formal budget from the board, he is willing to invest where a return on investment can be proven or if investment is an enabler for the strategic direction the company want to head in.

The CTO is also keen to improve the IT systems sufficiently so iSize are able to provide on demand prototypes at bicycle shows anywhere in the world in Pop Up event studios, the feasibility of this is currently being investigated.

Recently iSize were audited and it was reported that the network was implicitly open and communication was unobstructed from any device once access was gained internally. Data was also found in systems open to all employees and accessible by network admins within the Factory facility in Taiwan, this was stated as clear breach of the UK Data Protection Act 1998. A virus which created

a denial of Service to internal systems by continual broadcast was also recently introduced by a contract worker due to being unrestricted in using his own device on the network, in his defence he stated he had checked within HR documentation and could not locate any reason why he shouldn't use his own device and also had authorisation from his line manager. iSize has a firm stance on use of open protocols on all systems as they do not want to be tied to any specific vendor. iSize would like to run IPv6 at the earliest opportunity on the compute and network environments but have discovered the majority of the network infrastructure is not compatible and all applications will require extensive rework to provide IPv6 compatibility. For this reason IPv6 support has been suspended until a point in time when a clear return on investment can be proven.

Due to the UK recently voting to leave the EU and the resultant currency fluctuations marketing have submitted proposals to the board to commercialize on the opportunity of their products being effectively more cost effective as imports for the European market, as such they are keen to introduce five further sales offices in Paris, Berlin, Madrid, Rome & Zurich within the next year.

Document 2 - Network Background

iSize Head Quarters is located in Southampton, UK, it comprises of a Data Centre and user access network within the same facility for 320 users. Four sales sites in the UK (London. Oxford, Bristol & Ryde) connect to the HQ over a AT&T MPLS "wires only" network which transports voice (a central call manager for VoIP calls from branch to branch and branch to HQ calls is located within the DC with local PSTN breakout at each location) and data between sites. The MPLS network was originally provisioned in a Hub and Spoke topology with all inter branch traffic flowing through the Hub (Southampton). VoIP call setup flows through the call manager initially and then direct branch to branch voice traffic flows over the MPLS network, a voice pilot proved that if the voice traffic routed through the HQ between branch locations that the call quality suffered and employees regularly dialled off net to communicate with each other significantly raising PSTN call charges. Changing the Hub and Spoke MPLS topology into a standard any to any VPN rectified the call issues with VoIP traffic flowing direct between branches. Each sales site typically houses up to 10 employees with 10 VoIP handsets with a single homed customer owned CE router connected to a 48 port layer 2 switch, the switch has a separate voice and data VLAN which are trunked to the CE which then provides the default gateway functionality for each VLAN. Each VLAN is allocated a /24 subnet from the 10.11.0.0/16 address space with data VLANs being even and voice VLANs being odd in VLAN and subnet number allocation. A free Wi-Fi service is offered in each sales site for customers viewing product options, this is linked to local ADSL lines and typically dual Access Points connected on a unique VLAN on the layer 2 switch, the AP's provide DHCP services and a single SSID "iSize_bifi" which is unencrypted for easy access. Each sales site is connected to the MPLS network via 100Mbps Ethernet connections with a purchased access speed of 5Mbs. The HQ CE connection is 50Mbps over a 1Gbps circuit from a single customer owned CE dual attached to the collapsed core switches. QoS is currently not enabled on the WAN. All WAN locations access the internet via HQ to ensure traffic is filtered and monitored by the central firewalls by following a default route propagated to the MPLS network from the WAN CE in HQ (WANR1). The MPLS network also provides connectivity to the Bi-kea factory in Taiwan which connects into the MPLS network locally in Taiwan over a full rate DS3 connection, CAD designs and production plans are frequently downloaded from the DC to the factory.

The DC network is comprised of dual collapsed core modular switches and multiple layer 2 modular access switches, each switch has dual supervisors and the core switches have been fitted with dual Power Supply Units (PSUs), each core PSU is fed from a separate Power Distribution Unit (PDU). VRRP functionality for user and DC VLANs is provided by the layer 3 collapsed core central switches, CORESW1 is VRRP active for all VLANs and has been assigned a low root bridge ID to force root bridge selection, there are approx. 30% of ports spare but syslog messages have been seen stating the MAC table size is at 95% capacity . Users connect to switches ACCSW1 & ACCSW2 and all servers are single homed to the remaining access switches, each server access switch has a capacity of 240 1Gbps ports with approximately 50 free ports on each access switch, all card slots are in use in each switch. VLAN details are as follows:

VLAN Number	VLAN Use
1-10	Production VLANs
11-20	Test VLANs
21-30	User Access VLANs

VLANs are addressed from the 10.10.0.0/16 subnet with the 3rd octet matching the VLAN number. All VLANs are available and configured in each access switch to minimise any changes required when

new devices or users are added to the network. Changes typically take one week to perform and are chargeable via the 3rd party support company.

iSize has been assigned an AS of AS23 from their ISP (Fishnet, a Tier 2 ISP with AS55), their edge router EDGER1 runs eBGP to the ISP over a full 100Mbps Ethernet circuit and advertises their PA address space from the BGP enabled Firewalls FW1 & FW2 received over iBGP connections. The connections from the firewalls are active / standby with no state synchronization enabled between firewalls, synchronization and dynamic failover is achievable but only with a license upgrade. Active / Standby is currently achieved by FW1 advertising the assigned /24 public IPv4 range as two /25 prefixes to the EDGER1 router and by EDGER1 advertising a default route towards each firewall over BGP on two publicly registered /29 links. The firewalls also run OSPF and advertise default routes redistributed from BGP towards the core switches using Type 1 metrics with the links assigned a manual OSPF cost to ensure the default route advertised from FW1 is preferred for egress routing towards the internet. iSize employ a “Black List” approach to internet bound traffic.

A DMZ exists on the firewalls with two 48 port Layer 2 switches connected on a single VLAN, default gateway functionality for the DMZ hosts is provided by the firewalls using VRRP, FW1 is the VRRP active FW. The DMZ hosts public facing web services, DNS, NTP & Mail systems.

The network is managed via SNMPv2 polling and configuration management from a management server which is monitored by the in house Applications Help Desk Team.

The current network infrastructure is not capable of supporting IPv6. Fishnet (the ISP) do not currently offer IPv6 services on their network.

Network Diagram

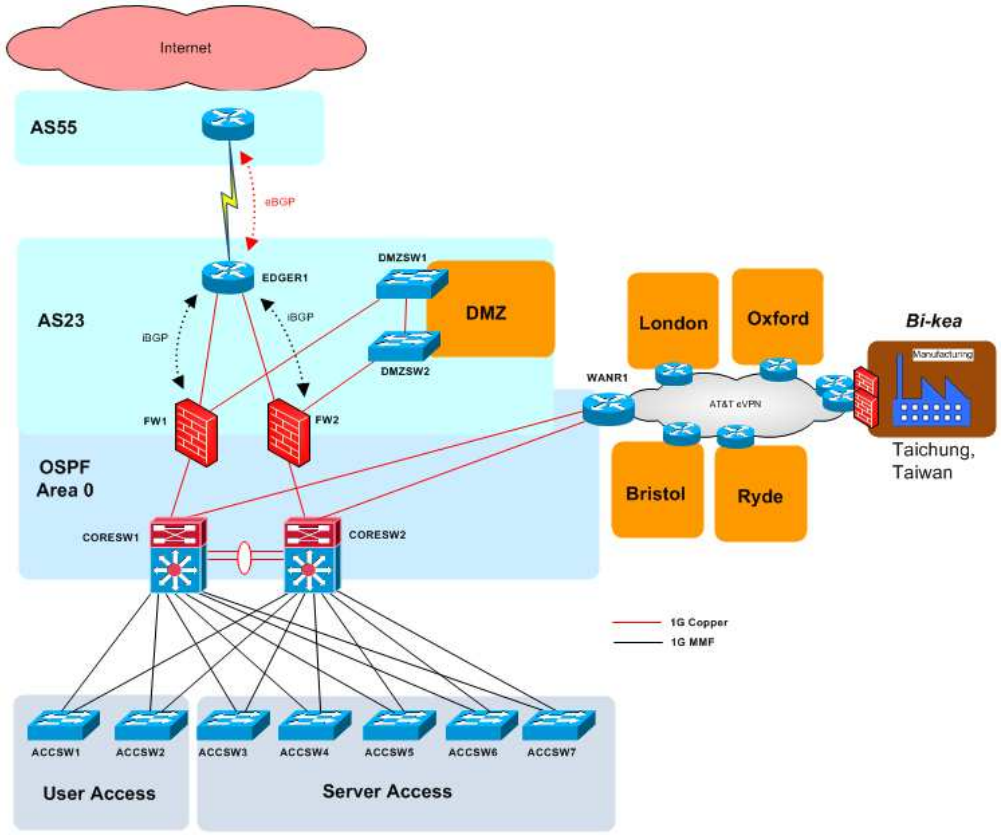

Q1) What are the six key technical issues facing iSize currently?

Issues	Key Technical Issues
VLAN1 is being used as a Production VLAN	
High number of unused / blocked links	
No onsite support or design authority	
No full mesh of iBGP between FW1 & FW2	
Single points of failure within the network	
Scalability	
Network Management	
High Opex	
Accessibility of shared systems from Taiwan	
No Security Policy present	
User and server access networks are not isolated with layer 3	
Breach of Data Protection Act	

Status %	6

Q2) What are the six key business issues facing iSize currently?

Issues	Key Business Issues
High Opex	
No onsite support or design authority	
Business Model Innovation	
Business Growth	
Lack of Scalability in the network	
Accessibility from Taiwan	
No Security Policy present	
Breach of Data Protection	
iSize could be liable for any maluses of internet access from Sales sites	

Status %	13

Q3) There is a requirement to add 200 new 1Gbps servers to the DC network. How would you recommend that iSize address the known scalability issue in the network to provide connectivity to the new server ports without replacing the existing collapsed core switches? Choose 1

a) Purchase additional line cards for the modular access switch chassis
b) Replace line cards with newer version cards with a higher density of aggregation ports
c) Replace 1Gbps line cards with 10Gbps line cards in the core
d) Purchase a small switch pair to act as an aggregation point for the user estate and then connect this to the existing collapsed core at layer 3
e) Purchase a small switch pair to act as aggregation point for the user estate and then connect this to the existing collapsed core at layer 2

Status %	**15**

Q4) iSize are keen to enhance resilience in the Access Layer of the DC, what would be the quickest and best return on investment they could make? Choose 1

a) Configure the access switches for VLAN load balancing towards the collapsed core to utilise unused trunk ports.
b) Add secondary supervisors to the access switches
c) Add additional PSUs to the access switches
d) Provide new NICs on each server and dual home them to different access switches with NIC teaming configured.
e) Replace the Access switches with new more reliable products

Status %	17

Q5) What would you recommend to iSize to strategically address the broadcast issue which arose from the virus introduced by the BYOD device? Choose 2

a) Introduce a security policy to include treatment of BYOD
b) Run broadcast control / suppression on trunk links
c) Run 802.1X Network Access Control paired with Radius & Active Directory on the user Access Switches to validate user credentials.

d) Separate the access layer and server layers with layer 3

Status %	19

Email1

From: Sam Dukes
To: CCDE Candidate
Subject: Thanks for your Observations

Hi

Thanks for highlighting our technical and business issues, I'm going to work on the business issues but I need you to sort out another issue just reported to me before we tackle the technical issues.
One of my server guys was investigating an issue where he felt he wasn't getting the performance he should be from his application, he was expecting a high bandwidth from ACCSW6 out to the sales offices for a new system on VLAN10 he span up but it was running like a dial up, he knows just enough about the switches to enable a session which copied the traffic on the trunk ports to a sniffer he setup on his laptop on the same switch. He found a huge amount of traffic destined to VLAN7 on one of the trunks of ACCSW6 and believes this is the root cause starving the available bandwidth to this switch and his application server, can you recommend a quick fix for this as our sales guys are complaining it's taking ages to access data in front of customers? We recently configured some VLAN load balancing so the core switches are VRRP active for odd and even VLANs but I don't think this has anything to do with it.

Rgds

Sam Dukes
CTO iSize